

EXPRESS MAIL LABEL NO.: EV 347799635	DATE OF DEPOSIT: <u>August 6, 2003</u>
I hereby certify that this paper and fee are being deposited with the United States Postal Service Express Mail Post Office to Addressee service under 37 CFR § 1.10 on the date indicated below and is addressed to Mail Stop Patent Applications, Commissioner for Patents, PO Box 1450, Alexandria, VA 22313-1450.	
<u>Sandra M. Campbell</u>	<u>Sandra M. Campbell</u>
NAME OF PERSON MAILING PAPER AND FEE	SIGNATURE OF PERSON MAILING PAPER AND FEE

Inventor(s): David Louis Kaminsky
David M. Ogle

AUTONOMIC ASSISTANCE FOR POLICY GENERATION

BACKGROUND OF THE INVENTION

Statement of the Technical Field

[0001] The present invention relates to the field of autonomic computing and more particularly to policy based autonomic management of a system.

Description of the Related Art

[0002] The task of modern systems administration differs significantly from that of days gone by. Not just a decade ago, systems administration primarily entailed the addition and deletion of users within the system, the management of print queues, disk maintenance, and the supervision and operation of daily backup procedures. Most if not all resources required by applications remained present in a single network, and few if any applications depended upon the operation of other, co-executing applications. In fact, the notion of an enterprise application, as compared to a mere network application remained largely within the realm of academia as a decade ago, the enabling technologies had not advanced enough in terms of speed and reliability to facilitate true enterprise computing.

[0003] Much has changed since the early days of network computing. Today, enterprise computing permeates the electronic landscape. While some enterprise

applications remain largely stand-alone, most rely in some respect on a co-existing enterprise application or a soft enterprise resource, such as a database application, Web application server, or other cooperating component. Thus, systems administration has advanced far beyond user and print queue administration and daily backup routines. Today, the interdependencies among system components present a significant challenge to the administrator. In this regard, the management of a single component in the system can depend upon the state of a multiplicity of other components. Thus, in administering the system, a seemingly endless array of administrative choices abound.

[0004] The task of systems administration recently has grown to include policy based management principles. Policy based management principles initially included rules for authentication only. Specifically, the rules specified which users defined within the system were permitted to perform which administrative tasks upon which components. Typically, the authority to perform such administrative tasks represents the sole type of rule managed by policy within the enterprise. Yet, more recent policy based management principles relate more specifically to the automated monitoring and maintenance of the system. In particular, policy based management has formed the cornerstone of "autonomic" systems administration.

[0005] In the famed manifesto, *Autonomic Computing: IBM's Perspective on the State of Information Technology*, Paul Horn, Senior Vice President of IBM Research, observed, "It's not about keeping pace with Moore's Law, but rather dealing with the consequences of its decades-long reign." Given this observation, Horn suggested a computing parallel to the autonomic nervous system of the biological sciences. Namely,

whereas the autonomic nervous system of a human being monitors, regulates, repairs and responds to changing conditions without any conscious effort on the part of the human being, in an autonomic computing system, the system must self-regulate, self-repair and respond to changing conditions, without requiring any conscious effort on the part of the computing system operator.

[0006] Thus, while the autonomic nervous system can relieve the human being from the burden of coping with complexity, so too can an autonomic computing system. Rather, the computing system itself can bear the responsibility of coping with its own complexity. The crux of the IBM manifesto relates to eight principal characteristics of an autonomic computing system:

- I. The system must "know itself" and include those system components which also possess a system identify.
- II. The system must be able to configure and reconfigure itself under varying and unpredictable conditions.
- III. The system must never settle for the status quo and the system must always look for ways to optimize its workings.
- IV. The system must be self-healing and capable of recovering from routine and extraordinary events that might cause some of its parts to malfunction.
- V. The system must be an expert in self-protection.
- VI. The system must know its environment and the context surrounding its activity, and act accordingly.
- VII. The system must adhere to open standards.

VIII. The system must anticipate the optimized resources needed while keeping its complexity hidden from the user.

[0007] Importantly, in accordance with the eight tenants of autonomic computing, several single system and peer-to-peer systems have been proposed in which self-configuration, management and healing have provided a foundation for autonomic operation. Self-managing systems which comport with the principles of autonomic computing reduce the cost of owning and operating computing systems. Yet, implementing a purely autonomic system has proven revolutionary. Rather, as best expressed in the IBM Corporation white paper, Autonomic Computing Concepts (IBM Corporation 2001)(hereinafter, the "IBM White Paper"), "Delivering system-wide autonomic environments is an evolutionary process enabled by technology, but it is ultimately implemented by each enterprise through the adoption of these technologies and supporting processes."

[0008] In the IBM White Paper, five levels have been logically identified for the path to autonomic computing. These five levels range from the most basic, manual process to the most purely autonomic. The Basic Level of autonomic computing represents a starting point of information technology environments. Each infrastructure element can be managed independently by an administrator who can establish, configure, monitor and ultimately replace the element. At the Managed Level of autonomic computing, systems management technologies can be used to collect information from disparate systems onto fewer consoles, reducing the time consumed for the administrator to collect and synthesize information as the environment becomes more complex.

[0009] Notably, the Predictive Level incorporates new technologies to provide a correlation among several infrastructure elements. These infrastructure elements can begin to recognize patterns, predict the optimal configuration of the system, and provide advice as to the nature of the course of action which the administrator ought to take. By comparison, at the Adaptive Level the system itself can automatically perform appropriate actions responsive to the information collected by the system and the knowledge of the state of the system. Finally, at the Autonomic Level the entire information technology infrastructure operation is governed by business policies and objectives. Users interact with the autonomic technology only to monitor the business processes, alter the objects, or both.

[0010] In the Predictive Level of autonomic systems administration, a system administration component can monitor the courses of action performed by a systems administrator in response to specific stimuli. Based upon the courses of action performed by the systems administrator, the system administration component can formulate a policy for autonomically responding to the same stimuli when in an Adaptive or Autonomic mode. The assumption that the systems administrator has performed appropriately responsive to the stimuli, however, may not always be a correct assumption. In many cases, the systems administrator lacks the experience to select the most optimal course of action. Moreover, in many cases, the systems administrator ought to consult with a more knowledgeable administrator in regard to the specific stimuli. Yet, often consultations are the exception rather than the rule. Thus, in the quest for more automation for efficiency, efficiency can be lost due to improper human decision making.

SUMMARY OF THE INVENTION

[0011] The present invention addresses the deficiencies of the art in respect to autonomic systems administration and provides a novel and non-obvious method, system and apparatus for autonomically assisting in the generation of a systems administration policy. A system for autonomically assisting in the creation of an administrative policy can include a systems administration component coupled to a system under study. A workflow component can be configured for communicative linkage to one or more policy makers. Specifically, the workflow component can include a further configuration for routing stimuli and response data from the system under study to a selected one of the policy makers based upon an affiliation between the selected policy maker and a portion of the system under study. Finally, a policy generation component can be coupled to the workflow component and configured to generate an administrative policy for administering the system under study based upon data collected from the selected one of the policy makers for the stimuli and response data.

[0012] A method for autonomically assisting in the creation of an administrative policy can include detecting a stimuli in a system under study and monitoring a response by a systems administrator to the stimuli. The stimuli and the response can be forwarded to a policy maker suited to analyze the stimuli and the response. Also, the policy maker can be queried for a preferred response to the stimuli. Based upon the preferred response to the stimuli, a policy for responding to the stimuli can be formulated. As such, the policy can be enforced in managing the system under study.

Moreover, the policy can be forwarded for future consideration to the systems administrator.

[0013] Additional aspects of the invention will be set forth in part in the description which follows, and in part will be obvious from the description, or may be learned by practice of the invention. The aspects of the invention will be realized and attained by means of the elements and combinations particularly pointed out in the appended claims. It is to be understood that both the foregoing general description and the following detailed description are exemplary and explanatory only and are not restrictive of the invention, as claimed.

BRIEF DESCRIPTION OF THE DRAWINGS

[0014] The accompanying drawings, which are incorporated in and constitute part of this specification, illustrate embodiments of the invention and together with the description, serve to explain the principles of the invention. The embodiments illustrated herein are presently preferred, it being understood, however, that the invention is not limited to the precise arrangements and instrumentalities shown, wherein:

[0015] Figure 1 is block diagram illustrating a system for autonomically assisting in the generation of a policy for use when responding to a particular stimulus;

[0016] Figure 2A is a flow chart illustrating a process for referring stimuli/response data to a workflow component in the system of Figure 1; and,

[0017] Figure 2B is a flow chart illustrating a process for generating a policy responsive to stimuli/response data received in the system of Figure 1.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

[0018] The present invention is a system, method and apparatus for autonomically assisting in the generation of a policy for use in the autonomic administration of a system. In accordance with the present invention, a systems administration component can monitor both stimuli generated in the system under study, and the response to the stimuli by a responding systems administrator. By stimuli it is meant that a particular situation has arisen giving rise to an uncertainty within the system in respect to how to proceed. A typical situation might include "disk is full", "backup scheduled" or "shutdown of database requested". In essence, a situation can arise when a state is encountered in the system which requires administration by the systems administrator. Similarly, by response it is meant that a particular course of action is chosen, such as "empty trash", "conduct backup", or "backup database before honoring shutdown request". Again, a response can include any course of action taken by the systems administrator in response to the stimuli.

[0019] Both the stimuli and the response can be stored for subsequent reference by a workflow process. In the workflow process, the stimuli and response can be routed to a qualified policy maker who can be queried for a suitable response in addition to other supplementary information. Other supplementary information can include, but is not limited to, related stimuli, related responses, and related systems in which the stimuli might arise. Based upon the data provided by the policy maker in response to the query, a policy can be formulated which can include a set of rules for responding to the stimuli. Subsequently, the policy can be established for use in an autonomic portion

of the systems administration component. Additionally, the policy can be provided to the systems administrator as well as the policy maker.

[0020] Figure 1 is block diagram illustrating a system for autonomically assisting in the generation of a policy for use when responding to a particular stimulus. The system of the present invention can include a systems administration component 115 coupled to a system under study 110. The systems administration component 115 can be configured in a number of modes which can include a predictive mode 120A, an adaptive mode 120B and an autonomic mode 120C. Other modes can be included such as a basic mode or a managed mode (neither shown), though the invention is not limited to any particular defined mode. Rather, as it will be apparent to the skilled artisan, the modal operation of the systems administration component 115 can be autonomic in respect to the formulation of a policy which can be utilized autonomically by the systems administration component 115, or manually by a systems administrator 105.

[0021] A policy generation process 125 can be coupled to both a data store of stimuli/responses 130 and a workflow process 135. The data store of stimuli/responses 130 can include a collection of stimuli flowing from the operation of the system 110, and the responses performed by the systems administrator 105. The workflow process 135 can be configured to route selected ones of the stimuli and responses to one or more policy makers 140A, 140B, 140n. The policy makers 140A, 140B, 140n can be administrators skilled in particular aspects of systems administration, business personnel qualified to analyze economic tradeoffs in selecting among suitable alternative courses of action, or some combination thereof. A table of policy makers

and particular stimuli (not shown) can be established so that the workflow process 135 can route workflow requests relating to particular stimuli to suitable ones of the policy makers 140A, 140B, 140n. Alternatively, the situation can be routed to a person (not shown) qualified to make the routing decision.

[0022] Based upon data provided by the policy makers 140A, 140B, 140n in analyzing routed ones of the workflow requests, the policy generation process 125 can produce a policy 145 having one or more rules 150A, 150B, 150n for handling the particular stimuli at hand. Specifically, the rules 150A, 150B, 150n can specify not only a preferred response to the stimuli, but also the rules can account for related stimuli and can include related actions. Finally, the policy 145 can be earmarked not only for the system under study 110, but also the policy 145 can be earmarked for use with other systems as specified by the policy makers 140A, 140B, 140n. Ordinarily, once formulated, the policy 145 can be shared with the systems administrator 105 and the policy makers 140A, 140B, 140n. Moreover, the policy 145 can be utilized by the systems administration component 115 when placed in a more autonomic mode.

[0023] Figure 2A is a flow chart illustrating a process for referring stimuli/response data to a workflow component in the system of Figure 1. Beginning in block 210, a stimuli can be detected in the system under study, and a response to the stimuli by a systems administrator further can be monitored in block 220. In block 230, both the stimuli and the response can be stored for reference by a policy generation component in Figure 2B. Finally, in block 240 a notification of the stimuli/response can be forwarded to a workflow component for routing to a suitable policy maker in which a policy can be generated for the stimuli.

[0024] Figure 2B is a flow chart illustrating a process for generating a policy responsive to stimuli/response data received in the system of Figure 1. Beginning in block 250, the notification of the stimuli/response can be received. In block 260, a suitable policy maker can be selected for the stimuli at hand. Specifically, the workflow component can refer to a pre-configuration in which certain policy makers have been associated with at least a particular stimuli, a particular system, a particular type of response, or a particular portion of a system. Once identified, in block 270 the stimuli/response can be routed to the selected policy maker.

[0025] In this regard, in blocks 280A through block 280D, the selected policy maker can be queried on any number of levels which can range from identifying a preferred response to the stimuli (block 280A) through identifying related actions, systems, and stimuli. Preferably, the selected policy maker can be queried through a series of dialog box frames, though the invention is not limited to the precise arrangement for presenting the queries to the selected policy maker. In any event, the response to the queries can be formulated into a policy in block 290 and in block 300 the policy can be distributed to the appropriate systems for enforcement. Subsequently, the resulting performance of the policy can be monitored and reported not only to the systems administrator, but also to the selected policy maker.

[0026] The present invention can be realized in hardware, software, or a combination of hardware and software. An implementation of the method and system of the present invention can be realized in a centralized fashion in one computer system, or in a distributed fashion where different elements are spread across several interconnected computer systems. Any kind of computer system, or other apparatus

adapted for carrying out the methods described herein, is suited to perform the functions described herein.

[0027] A typical combination of hardware and software could be a general purpose computer system with a computer program that, when being loaded and executed, controls the computer system such that it carries out the methods described herein. The present invention can also be embedded in a computer program product, which comprises all the features enabling the implementation of the methods described herein, and which, when loaded in a computer system is able to carry out these methods.

[0028] Computer program or application in the present context means any expression, in any language, code or notation, of a set of instructions intended to cause a system having an information processing capability to perform a particular function either directly or after either or both of the following a) conversion to another language, code or notation; b) reproduction in a different material form. Significantly, this invention can be embodied in other specific forms without departing from the spirit or essential attributes thereof, and accordingly, reference should be had to the following claims, rather than to the foregoing specification, as indicating the scope of the invention.